

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments
1	BRS	L1	0	exception SAME interrupt SAME multiprocessor SAME exception adj mode	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 13:48	
2	BRS	L4	0	interrupt SAME multiprocessor SAME exception adj mode	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 13:49	
3	BRS	L5	0	multiprocessor SAME exception adj mode	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 13:48	
4	BRS	L6	72	interrupt SAME multiprocessor SAME exception	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 13:49	
5	BRS	L7	23	6 and register WITH mode	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 13:49	
6	BRS	L8	12	7 and exception adj hand\$5	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:06	
7	BRS	L9	3	8 and exception adj vector adj table	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:09	

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments
8	BRS	L10	68	exception adj vector adj table	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:09	
9	BRS	L11	26	10 and (shared near3 memory)	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:15	
10	BRS	L12	25	11 not 9	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:10	
11	BRS	L13	25	11 not 8	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:10	
12	BRS	L14	13	13 and multiprocessor	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:19	
13	BRS	L15	1	(exception adj vector adj table) WITH (shared near3 memory)	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:16	
14	BRS	L16	2	(exception adj vector) WITH (shared near3 memory)	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:19	

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments
15	BRS	L17	71	(exception adj vector) and (shared near3 memory)	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:19	
16	BRS	L18	43	17 and multiprocessor	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:19	
17	BRS	L19	43	18 and register	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:19	
18	BRS	L20	0	19 and interrupt adj request	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:20	
19	BRS	L21	6	19 and interrupt adj request	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:20	
20	BRS	L22	6	21 not 14	USPAT; US-PGP UB; EPO; JPO; DERWE NT; IBM_TD B	2004/09/01 14:20	


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

multiprocessor and exception and "vector table" and interrupt



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used **multiprocessor** and **exception** and **vector table** and **interrupt** and **mode** and **RISC** and **shared memory** and **register** and **interrupt request**

Found 6,373 of 141,680

Sort results by

relevance

Display results

expanded form

☒ [Save results to a Binder](#)
☒ [Search Tips](#)
☐ [Open results in a new window](#)

 Try an [Advanced Search](#)

 Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Hardware and software support for efficient exception handling](#)

Chandramohan A. Thekkath, Henry M. Levy

 November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29, 28 Issue 11, 5

Full text available: pdf (1.44 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Program-synchronous exceptions, for example, breakpoints, watchpoints, illegal opcodes, and memory access violations, provide information about exceptional conditions, interrupting the program and vectoring to an operating system handler. Over the last decade, however, programs and run-time systems have increasingly employed these mechanisms as a performance optimization to detect normal and expected conditions. Unfortunately, current archi ...

2 [Decoupled hardware support for distributed shared memory](#)

Steven K. Reinhardt, Robert W. Pfile, David A. Wood

 May 1996 **ACM SIGARCH Computer Architecture News, Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf (1.47 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper investigates hardware support for fine-grain distributed shared memory (DSM) in networks of workstations. To reduce design time and implementation cost relative to dedicated DSM systems, we decouple the functional hardware components of DSM support, allowing greater use of off-the-shelf devices. We present two decoupled systems, Typhoon-0 and Typhoon-1. Typhoon-0 uses an off-the-shelf protocol processor and network interface; a custom access control device is the only DSM-specific hard ...

3 [An analysis of operating system behavior on a simultaneous multithreaded architecture](#)

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy

 November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28, 34 Issue 5, 5

Full text available: pdf (227.80 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


This paper presents the first analysis of operating system execution on a simultaneous multithreaded (SMT) processor. While SMT has been studied extensively over the past 6 years, previous research has focused entirely on user-mode execution. However, many of the applications most amenable to multithreading technologies spend a significant fraction of

their time in kernel code. A full understanding of the behavior of such workloads therefore requires execution and measurement of the operating sy ...

4 An analysis of operating system behavior on a simultaneous multithreaded architecture

Joshua A. Redstone, Susan J. Eggers, Henry M. Levy

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available:  pdf (1.56 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents the first analysis of operating system execution on a simultaneous multithreaded (SMT) processor. While SMT has been studied extensively over the past 6 years, previous research has focused entirely on user-mode execution. However, many of the applications most amenable to multithreading technologies spend a significant fraction of their time in kernel code. A full understanding of the behavior of such workloads therefore requires execution and measurement of the operating sy ...

5 Software-controlled caches in the VMP multiprocessor

D. R. Cheriton, G. A. Slavenburg, P. D. Boyle

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available:  pdf (958.63 KB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

VMP is an experimental multiprocessor that follows the familiar basic design of multiple processors, each with a cache, connected by a shared bus to global memory. Each processor has a synchronous, virtually addressed, single master connection to its cache, providing very high memory bandwidth. An unusually large cache page size and fast sequential memory copy hardware make it feasible for cache misses to be handled in software, analogously to the handling of virtual memory page faults. Har ...

6 A VLIW architecture for a trace scheduling compiler

Robert P. Colwell, Robert P. Nix, John J. O'Donnell, David B. Papworth, Paul K. Rodman

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available:  pdf (1.59 MB)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Very Long Instruction Word (VLIW) architectures were promised to deliver far more than the factor of two or three that current architectures achieve from overlapped execution. Using a new type of compiler which compacts ordinary sequential code into long instruction words, a VLIW machine was expected to provide from ten to thirty times the performance of a more conventional machine built of the same implementation technology. Multiflow Computer, Inc., has now built a VLIW called the TRACE™ < ...

7 Automatic Generation of Fast Timed Simulation Models for Operating Systems in SoC Design

S. Yoo, G. Nicolescu, L. Gauthier, A. Jerraya

March 2002 **Proceedings of the conference on Design, automation and test in Europe**

Full text available:  pdf (158.87 KB)

Additional Information: [full citation](#), [abstract](#)

 [Publisher Site](#)


To enable fast and accurate evaluation of HW/SW implementation choices of on-chip communication, we present a method to automatically generate timed OS simulation models. The method generates the OS simulation models with the simulation environment as a virtual processor. Since the generated OS simulation models use final OS code, the presented method can mitigate the OS code equivalence problem. The generated model also simulates different types of processor exceptions. This approach provides two orders ...

8 OMP: a RISC-based multiprocessor using orthogonal-access memories and multiple spanning buses

K. Hwang, M. Dubois, D. K. Panda, S. Rao, S. Shang, A. Uresin, W. Mao, H. Nair, M. Lytwyn, F.

Hsieh, J. Liu, S. Mehrotra, C. M. Cheng

June 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 4th international conference on Supercomputing**, Volume 18 Issue 3

Full text available:  pdf(1.96 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the architectural design and RISC based implementation of a prototype supercomputer, namely the Orthogonal MultiProcessor (OMP). The OMP system is constructed with 16 Intel 1860 RISC microprocessors and 256 parallel memory modules, which are 2-D interleaved and orthogonally accessed using custom-designed spanning buses. The architectural design has been validated by a CSIM-based multiprocessor simulator. The design choices are based on worst-case delay a ...

9 Making operating systems more robust: Improving the reliability of commodity operating systems



Michael M. Swift, Brian N. Bershad, Henry M. Levy

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available:  pdf(262.78 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite decades of research in extensible operating system technology, extensions such as device drivers remain a significant cause of system failures. In Windows XP, for example, drivers account for 85% of recently reported failures. This paper describes Nooks, a *reliability subsystem* that seeks to greatly enhance OS reliability by isolating the OS from driver failures. The Nooks approach is practical: rather than guaranteeing complete fault tolerance through a new (and incompatible) OS ...

Keywords: I/O, device drivers, protection, recovery, virtual memory

10 Automatic generation of application-specific architectures for heterogeneous multiprocessor system-on-chip



Damien Lyonnard, Sungjoo Yoo, Amer Baghdadi, Ahmed A. Jerraya

June 2001 **Proceedings of the 38th conference on Design automation**

Full text available:  pdf(285.15 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


We present a design flow for the generation of application-specific multiprocessor architectures. In the flow, architectural parameters are first extracted from a high-level system specification. Parameters are used to instantiate architectural components, such as processors, memory modules and communication networks. The flow includes the automatic generation of communication coprocessor that adapts the processor to the communication network in an application-specific way. Experiments with ...

11 The interaction of architecture and operating system design



Thomas E. Anderson, Henry M. Levy, Brian N. Bershad, Edward D. Lazowska

April 1991 **Proceedings of the fourth international conference on Architectural support for programming languages and operating systems**, Volume 26 , 19 , 25 Issue 4 , 2 , Special Issue

Full text available:  pdf(1.60 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 Cellular disco: resource management using virtual clusters on shared-memory multiprocessors



Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Full text available:  pdf(287.05 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource management, scalable multiprocessors, virtual machines

13 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available:  [pdf \(1.93 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a system ...

14 The Clipper processor: instruction set architecture and implementation

W. Hollingsworth, H. Sachs, A. J. Smith

February 1989 **Communications of the ACM**, Volume 32 Issue 2

Full text available:  [pdf \(4.67 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Intergraph's CLIPPER microprocessor is a high performance, three chip module that implements a new instruction set architecture designed for convenient programmability, broad functionality, and easy future expansion.

15 Poster session 2: Cycle-accurate power analysis for multiprocessor systems-on-a-chip

Mirko Loghi, Massimo Poncino, Luca Benini

April 2004 **Proceedings of the 14th ACM Great Lakes symposium on VLSI**

Full text available:  [pdf \(191.54 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Developing energy-aware software for multiprocessor systems-on-chip (MPSoCs) is a difficult task, which requires the knowledge of the distribution of the power consumption among several heterogeneous devices (cores, memories, busses, etc.). In this work we analyze the power breakdowns of power consumption for a complete MPSoC platform, under several application workloads and operating conditions. We leverage a complete-system simulation platform with accurate power models for all key hardware mo ...

Keywords: low power, multiprocessor, system-on-chip

16 Integration of message passing and shared memory in the Stanford FLASH multiprocessor

John Heinlein, Kourosh Gharachorloo, Scott Dresser, Anoop Gupta

November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29, 28 Issue 11, 5

Full text available:  [pdf \(1.80 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The advantages of using message passing over shared memory for certain types of communication and synchronization have provided an incentive to integrate both models within a single architecture. A key goal of the FLASH (FLexible Architecture for SHared memory) project at Stanford is to achieve this integration while maintaining a simple and efficient design. This paper presents the hardware and software mechanisms in FLASH to support various message passing protocols. We achieve low overhead ...

17 Embrea: fast and flexible machine simulation

Emmett Witchel, Mendel Rosenblum

May 1996 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1996 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 24 Issue 1

Full text available:  pdf(1.83 MB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes Embrea, a simulator for the processors, caches, and memory systems of uniprocessors and cache-coherent multiprocessors. When running as part of the SimOS simulation environment, Embrea models the processors of a MIPS R3000/R4000 machine faithfully enough to run a commercial operating system and arbitrary user applications. To achieve high simulation speed, Embrea uses dynamic binary translation to generate code sequences which simulate the workload. It is the first machine simu ...

18 Hardware fault containment in scalable shared-memory multiprocessors

Dan Teodosiu, Joel Baxter, Kinshuk Govil, John Chapin, Mendel Rosenblum, Mark Horowitz

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2

Full text available:  pdf(2.05 MB)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Current shared-memory multiprocessors are inherently vulnerable to faults: any significant hardware or system software fault causes the entire system to fail. Unless provisions are made to limit the impact of faults, users will perceive a decrease in reliability when they entrust their applications to larger machines. This paper shows that fault containment techniques can be effectively applied to scalable shared-memory multiprocessors to reduce the reliability problems created by increased mach ...

19 Alpha AXP architecture

Richard L. Sites

February 1993 **Communications of the ACM**, Volume 36 Issue 2

Full text available:  pdf(4.62 MB)


Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

Keywords: Alpha AXP chip

20 VLSI assist for a multiprocessor

Bob Beck, Bob Kasten, Shreekanth Thakkar

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available:  pdf(1.28 MB)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Multiprocessors have long been of interest to computer community. They provide the potential for accelerating applications through parallelism and increased throughput for large multi-user system. Three factors have limited the commercial success of multiprocessor systems; entry cost, range of performance, and ease of application. Advances in very large scale integration (VLSI) and in computer aided design (CAD) have removed these limitations, making possible a new class of multiprocessor system ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [!\[\]\(d66ff64371a51729ac8c1cdaa685ba6f_img.jpg\) Adobe Acrobat](#) [!\[\]\(0f31ebba7abcd47777e178db26f29705_img.jpg\) QuickTime](#) [!\[\]\(63ea948177b1bcc486b2b76d20d5fb69_img.jpg\) Windows Media Player](#) [!\[\]\(886f7dced1265a6d438eca0881817b40_img.jpg\) Real Player](#)